arXiv:1701.08140v1 [stat.ME] 27 Jan 2017

Network classification with applications to brain connectomics

Jesús D. Arroyo Relión, Daniel Kessler, Elizaveta Levina, Stephan F. Taylor
{jarroyor,kesslerd,elevina,sftaylor}@umich.edu

University of Michigan

January 30, 2017

Abstract

While statistical analysis of a single network has received a lot of attention in recent years, with a focus on social networks, analysis of a sample of networks presents its own challenges which require a different set of analytic tools. Here we study the problem of classification of networks with labeled nodes, motivated by applications in neuroimaging. Brain networks are constructed from imaging data to represent functional connectivity between regions of the brain, and previous work has shown the potential of such networks to distinguish between various brain disorders, giving rise to a network (or graph) classification problem. Existing approaches to graph classification tend to either treat all edge weights as a long vector, ignoring the network structure, or focus on the graph topology while ignoring the edge weights. Our goal here is to design a graph classification method that uses both the individual edge information and the network structure of the data in a computationally efficient way. We are also interested in obtaining a parsimonious and interpretable representation of differences in brain connectivity patterns between classes, which requires variable selection. We propose a graph classification method that uses edge weights as variables but incorporates the network nature of the data via penalties that promotes sparsity in the number of nodes. We implement the method via efficient convex optimization algorithms and show good performance on data from two fMRI studies of schizophrenia.

1 Introduction

Network data analysis has received a lot of attention in recent literature, especially unsupervised analysis of a single network which is thought of as generated from an exchangeable random graph model, for example Bickel and Chen (2009); Le et al. (2015); Zhang and Zhou (2015); Gao et al. (2015) and many others. This setting is applicable to a number of real life scenarios, such as social networks, but there are situations where the network nodes are labeled and therefore not exchangeable, and/or more than one network is available for analysis, which have received relatively less attention. Here we focus on the setting motivated by brain connectomics studies, where a sample of networks is available from multiple populations of interest (for example, mentally ill patients and healthy controls). In this setting, each unit in the population (e.g., a patient) is represented by their own network, and the nodes (brain regions of interest) are labeled and shared across all networks, through a registration process that maps all individual brains onto a common atlas. There are many classical statistical inference questions one can ask in this setting, for example, how to compare different populations (Tang et al., 2014a,b). The question we focus on in this paper is a classification

problem: given a training sample of networks with labeled nodes drawn from multiple classes, the goal is to learn the rules for predicting the class of a given network, and just as importantly, interpret these rules.

Network methods are a popular tool in the neuroscience literature, since brain networks have been shown to differentiate between patients with certain illnesses, or between types of external stimuli (see Bullmore and Sporns (2009) for a review). How brain connectivity networks are constructed varies with the type of imaging technology used, but in most cases connections cannot be observed directly, especially in humans. Here we focus on networks constructed from functional magnetic resonance imaging (fMRI) data, although the methods we develop are applicable to any sample of weighted networks with labeled nodes. In fMRI studies, BOLD (blood oxygen-level dependent) signal, a known correlate of underlying neural activity, is measured at a sequence of time points at many spatial locations, known as voxels, in the brain, resulting in a 4-dimensional data array, with three spatial dimensions and time. All the data in this paper come from resting state fMRI data, where patients are not asked to perform any particular task but are free to think about anything they want. The data arrays undergo pre-processing and registration, discussed in greater detail below, after which a connectivity measure is computed for every pair of nodes (Smith et al., 2013), most commonly the Pearson correlation coefficient between locations computed by averaging over the time dimension. It has been argued that partial correlations are a better measure of connectivity (Varoquaux and Craddock, 2013; Narayan et al., 2015), but the choice depends on the final goal of analysis, and in this paper we follow the vast majority of the connectomics literature and measure connectivity by marginal correlations between the corresponding time series. However, because of concerns about how global signal regression may alter the rank ordering of correlations, we instead directly use the ranks of the marginal correlation values (Gotts et al., 2013). In addition, we find that in our setting ranks of marginal correlations outperform other measures, such as partial correlations, in terms of classification error.

The general problem of graph classification has been studied previously in other contexts, with a substantial literature motivated by the problem of classification of chemical compounds (Srinivasan et al., 1996; Helma et al., 2001), where graphs represent the compound's molecular structure. Classification methods for chemical compounds are usually based in finding certain patterns in the graphs, like subgraphs or paths (Gonzalez et al., 2000; Kashima et al., 2003), that are more present in one class than the other. Since searching for patterns is a combinatorial task, the search is usually done in a greedy fashion or for a limited subgraph size. Once the differentiating patterns are found, a classifier can be trained using standard classification methods (Deshpande et al., 2005; Inokuchi et al., 2000; Kudo et al., 2004; Fei and Huan, 2010). In this setting the edges are binary or categorical (representing types of connections), whereas in brain networks the edges are naturally weighted since they are inferred from correlations. The molecular structure of the compounds is known so there is no noise, and typical graphs are very small, with about 20 nodes on average (Ketkar et al., 2009) which differ across networks, whereas brain networks are typically calculated over hundreds or thousands of nodes but all the networks share the same set of nodes. Thus, this class of methods is not suitable in our context.

In the context of brain networks classification, two main approaches have been followed. One way is to reduce a network to global summary measures such as the average degree, the clustering coefficient, or the average path length (Bullmore and Sporns, 2009), and use them as features. Hypothesis testing on a type of graph average has also been proposed (Ginestet et al., 2014). Although these features have been shown to be useful as diagnostic biomarkers for certain diseases (Supekar et al., 2008; Liu et al., 2008), global summary statistics do not allow for interpretation on a local scale, and miss network differences that manifest locally. One alternative is to compute these graph metrics for pre-defined regions and use them in combination with

a supervised learning method (Zhang et al., 2012), which is a compromise between between global summaries and vectorization approach, described next.

The vectorization approaches treats edge weights as a "bag of features" without taking into account the underlying network structure. These approaches have the advantage of leveraging the many existing high-dimensional classification techniques, and they provide an interpretation at the edge level if variable selection is applied (Richiardi et al., 2011; Craddock et al., 2009). Moreover, spatial correlation between edges connecting nodes that are close in the brain can be incorporated (Watanabe et al., 2014), although the effectiveness of this regularization will depend on the parcellation used to define the nodes (Power et al., 2011). Alternatively, an individual test can be used for each edge to find significant differences between the two populations, with a multiple testing correction and without constructing a classifier at all (Narayan et al., 2015). While vectorization methods can lead to good prediction performance, their interpretability is limited to individual edge selection, which is less scientifically interesting than identifying differentiating regions, and they cannot account for network structure. Taking the network structure into account can have benefits for both testing and classification settings. Other methods have been introduced that perform inference over groups of edges based on the community assignments of the nodes to which they are incident. For example, Sripada et al. (2014a,b) introduced *Network Contingency Analysis* which begins with massive univariate testing at each edge, and then counts the number of superthreshold connections in each *cell* (where a cell is defined as a group of edges that connection nodes in two functional systems). Nonparametric methods are then used to conduct inference on the count statistic for each cell, with multiple comparison correction for inference at the cell level. In multiple testing, power can be improved by applying a network-based dependence correction (Zalesky et al., 2010). For classification, better interpretability and potentially accuracy can be obtained if we focus on understanding which brain regions or interactions between them are responsible for the differences. Somewhat in this spirit, Vogelstein et al. (2013) proposed to look for a minimal set of nodes which best explains the difference, though that requires solving a combinatorial problem.

Our goal in this paper is to develop a high-dimensional network classifier that uses all the individual edge weights but also respects the network structure of the data and produces more interpretable results. To achieve this goal, we use structured sparsity penalties to incorporate the network information by penalizing both the number of edges and the number of nodes selected. Although our main application here is classification of brain connectivity networks, our methods are applicable to any weighted graphs with labeled nodes, and to general response variables rather than just class labels.

The rest of this paper is organized as follows. In Section 2, we introduce our classifier and the structured penalties. In Section 3 we show how to efficiently solve the resulting convex optimization problem by a proximal algorithm, each step of which is a further optimization problem which we solve by the alternating direction method of multipliers (ADMM). In Section 5, we analyze two brain connectivity datasets, each containing schizophrenic patients and healthy controls, and show that our regularization framework leads to state-of-the-art accuracy while providing interpretable results, some of which are consistent with previous findings and some are new. We conclude with a brief discussion in Section 6.

2 A framework for node selection in graph classification

2.1 A penalized graph classification approach

We start from setting up notation. All graphs we consider are defined on the same set of N labeled nodes. A graph can be represented with its adjacency matrix $A \in \mathbb{R}^{N\times N}$. We focus on graphs that are undirected ($A_{ij} = A_{ji}$) and contain no self-loops ($A_{ii} = 0$). These assumptions are not required for the derivations below, but they match the neuroimaging setting and simplify notation. Our goal is predicting a class label Y from the graph adjacency matrix A; in this paper we focus on the binary classification problem where Y takes values $\{-1, 1\}$, although extensions from binary to multi-class classification or real-valued responses are straightforward. Throughout this paper, we use $\|\cdot\|_p$ to denote the entry-wise ℓ_p norm, so for a matrix $A \in \mathbb{R}^{m\times n}$, $\|A\|_p = \left(\sum_{i=1}^m \sum_{j=1}^n |A_{ij}|^p\right)^{1/p}$.

A standard general approach is to construct a linear classifier, which predicts the response Y from a linear combination of the elements of A, $\langle A, B\rangle = \mathrm{Tr}\left(B^T A\right)$, where we arrange the coefficients in a matrix $B \in \mathbb{R}^{N\times N}$ to emphasize the network nature of the predictors. We focus on linear classifiers here because variable selection is at least as important as prediction itself in the neuroimaging application, and setting coefficients to 0 is a natural way to achieve this. The coefficients are typically estimated from training data by minimizing an objective consisting of a loss function plus a penalty. The penalty can be used to regularize the problem to make the estimator well-defined in high-dimensional problems, to select important predictors, and to impose structure, and many such penalties have been proposed, starting from the lasso (Tibshirani, 1996). Our focus in this work is on designing a classifier in this framework that respects and utilizes the network nature of the predictors. In brain networks in particular, neuroscientists believe that edges are organized in subnetworks, also called brain systems (Power et al., 2011), that carry out specific functions, and certain subnetworks are important for prediction (Bullmore and Sporns, 2009), although different studies tend to implicate different regions (Fornito et al., 2012). Thus we aim to find subnetworks with good discriminative power, and hence select both the most informative nodes and edges.

Let $\{(A^{(1)}, Y_1), \ldots, (A^{(n)}, Y_n)\}$ be the training sample of undirected adjacency matrices with their class labels, and let $Y = (Y_1, \ldots, Y_n)$. A generic linear classifier described above is computed by finding the coefficients B defined by

$$\hat{B} = \underset{B\in\mathcal{B}}{\arg\min} \left\{\ell(B) + \Omega(B)\right\}, \tag{2.1}$$

where $\mathcal{B} = \left\{B \in \mathbb{R}^{N\times N} : B = B^T, \mathrm{diag}(B) = 0\right\}$, Ω is a penalty, and

$$\ell(B) = \frac{1}{n}\sum_{k=1}^n \tilde{\ell}(Y_k, A^{(k)}; B)$$

is the training loss. Many choices of loss function $\tilde{\ell}$ are possible here, but in order to facilitate optimization, it is convenient to use a convex and continuously differentiable loss. Here we use the logistic loss function

$$\tilde{\ell}(Y, A; B, b) = \log\left(1 + \exp(-Y\langle A, B\rangle + b)\right),$$

including the threshold b as an additional parameter to be estimated.

To capture structural assumptions on important predictive edges, we focus on convex structured sparsity penalties (Bach et al., 2012) that encourage a small number of active nodes, by which we mean nodes attached to at least one edge with a non-zero coefficient. One approach to finding a set of such nodes was proposed

by Vogelstein et al. (2013), who called it a signal-subgraph, and proposed finding the minimal set of nodes (called signal vertices) which together are incident to all selected edges (but not every node connected to a selected edge is a signal vertex). Finding this set is a combinatorial optimization problem, and the set is not always uniquely defined. Instead, we focus on convex formulations that allow for efficient computation and encourage small active node sets indirectly.

Other convex penalties have been used for fMRI data as a way to enforce spatial smoothness in the solution (Grosenick et al., 2013; Xin et al., 2014; Hu and Allen, 2015). These methods assume that voxels are equally spaced in the brain, and neighboring voxels are highly correlated. In particular, Watanabe et al. (2014) proposed penalties for brain network classification using these spatial assumptions. Here, instead of enforcing a spatial regularization directly, we aim for a regularization that can be applied to any type of network data, and in particular to brain networks with coarse and/or uneven parcellations where enforcing spatial smoothness may not work as well. In any case, the flexibility of convex optimization algorithms allows one to easily incorporate additional spatially-informed penalties if needed.

2.2 Selecting nodes and edges through group lasso

To reflect the network structure of the predictors, we use a penalty that promotes a sparse classifier not only in the number of edges used, but also in the number of nodes. The group lasso penalty Yuan and Lin (2006) is designed to eliminate a group of variables simultaneously. Here we penalize the number of active nodes by treating all edges connected to one node as a group. Then eliminating this group (a row of coefficients in the matrix B) is equivalent to de-activating a node. The group penalty is defined as

$$\Omega_{\lambda,\rho}(B) = \lambda \left(\sum_{i=1}^{N} \|B_{(i)}\|_2 + \rho \|B\|_1 \right), \tag{2.2}$$

where $B_{(i)}$ is the i-th row of B (or equivalently, the i-th column), $\|B\|_1 = \sum_i \sum_j |B_{ij}|$ is the element-wise ℓ_1 norm of B and $\lambda, \rho \geq 0$ are tuning parameters. Note that the constraint $B = B^T$ makes the groups overlap, since a coefficient B_{ij} belongs to groups associated with the nodes i and j, and therefore, the edge between nodes i and j would be selected only if neither node was de-activated. The second term in the penalty $\rho\|B\|_1$ acts as the usual lasso penalty to promote sparsity inside the group (Friedman et al., 2010a), allowing to select a subset of edges for an active node. Due to the overlap in the groups, this lasso penalty is necessary in order to produce sparse solution. The constraint $\text{diag}(B) = 0$ in (2.1) is automatically enforced with this formulation.

Remark 1. An alternative to the constraint in the problem (2.1) is to optimize over the set

$$\tilde{\mathcal{B}} = \left\{ B \in \mathbb{R}^{N \times N}, \text{diag}(B) = 0 \right\}.$$

Without the symmetry constraint and assuming undirected graphs, the penalty (2.2) is equivalent to the overlapping group lasso formulation of Jacob et al. (2009). This formulation has some advantages. Since it gives group lasso without overlaps, the lasso penalty $\rho\|B\|_1$ is not required to obtain sparse solutions, and more efficient optimization algorithms exist for this case. This approach would loosely correspond to the idea of selecting signal nodes as in Vogelstein et al. (2013), in the sense that an edge can be selected if at least one of its nodes is selected, and the second node could be inactive. The downside is that each edge now corresponds to two different coefficients B_{ij} and B_{ji}, the problem encountered by all variable selection

methods that ignore symmetry, such as Meinshausen and Bühlmann (2006). The standard solution for this problem, as suggested by Jacob et al. (2009), is to take the average of the coefficients. Intuitively, one would expect that the formulation using $\mathcal{B}$ would be better when the significant edges are incident to a small set of nodes, since both nodes have to be active for an edge to be selected, while using $\tilde{\mathcal{B}}$ may be better when for some nodes most of their edges are significant, creating "significant hubs". Since in our application we are primarily looking for discriminative brain subnetworks, we focus on the symmetrically constrained formulation for the rest of the paper. We also found that in practice this second formulation results in less accurate classifiers for the neuroimaging data discussed in Section 5.

Remark 2. The analogue to (2.2) for directed graphs would assign coefficients B_{ij} and B_{ij} to the same group, resulting in the penalty

$$\Psi_{\lambda,\rho}(B) = \lambda \left(\sum_{i=1}^{N} \sqrt{\sum_{j} \left(B_{ij}^2 + B_{ji}^2 \right)} + \rho \|B\|_1 \right), \tag{2.3}$$

where $B \in \tilde{\mathcal{B}}$. Alternatively, we can also use the formulation of Remark 1, by replicating the variables and estimating two matrices of coefficients, say $B^{(1)}$ and $B^{(2)}$, with the penalty

$$\tilde{\Psi}_{\lambda,\rho}(B^{(1)}, B^{(2)}) = \lambda \left(\sum_{i=1}^{N} \sqrt{\sum_{j} \left(\left(B_{ij}^{(1)} \right)^2 + \left(B_{ji}^{(2)} \right)^2 \right)} + \rho \left(\|B^{(1)}\|_1 + \|B^{(2)}\|_1 \right) \right),$$

with $B^{(1)}, B^{(2)} \in \tilde{\mathcal{B}}$, and set the coefficients matrix to $B = \left(B^{(1)} + B^{(2)} \right)/2$. This formulation will again not directly select subnetworks as discussed in Remark 1.

Finally, for numerical stability we add an extra ridge penalty term $\frac{\gamma}{2}\|B\|_F^2 = \frac{\gamma}{2}\mathrm{Tr}\left(B^T B\right)$, with γ a small and fixed constant. This is common practice in high-dimensional prediction. Putting everything together, we solve the problem

$$(\hat{B}, \hat{b}) = \underset{B \in \mathcal{B}, b \in \mathbb{R}}{\arg\min} \left\{ \frac{1}{n} \sum_{k=1}^{n} \log \left(1 + \exp(-Y_k \langle B, A^{(k)} \rangle + b \right) + \frac{\gamma}{2} \|B\|_F^2 + \lambda \left(\sum_{i=1}^{N} \|B_{(i)}\|_2 + \rho \|B\|_1 \right) \right\} \tag{2.4}$$

for given values of λ and ρ, which will be chosen by cross-validation.

3 The optimization algorithm

Our optimization algorithm to solve the problem (2.4) combines two common approaches to convex optimization: proximal algorithms and alternating direction method of multipliers (ADMM). We use an accelerated version of the proximal algorithm (Beck and Teboulle, 2009) to solve the main problem (2.4). In each step, we need to calculate a proximal operator, which is a further convex optimization problem solved with the ADMM algorithm.

The main optimization difficulty comes from the overlapping groups. Some algorithms have been proposed for this case, including a subgradient descent method (Duchi and Singer, 2009), which has a slow rate of convergence, or a proximal algorithms based on smoothing the original problem (Yuan et al., 2011; Chen et al., 2012). Although smoothing yields fast algorithms, it is not clear that the sparsity pattern is preserved with those approximations. We follow an approach similar to Yuan et al. (2011) and Chen et al. (2012), but solve the proximal operator for the penalty (2.2) directly using the ADMM method. This can potentially

give a more accurate sparsity pattern, and the flexibility of the algorithm allows for additional penalties if desired, such as spatial smoothing similar to Watanabe et al. (2014) (see Remark 3).

The main problem (2.1) is solved with a proximal algorithm (see Parikh and Boyd (2013)). Recall that the proximal operator for a function f is defined as $\text{prox}_f(v) = \arg\min_x\{f(x) + \frac{1}{2}\|x - v\|_2^2\}$. Starting with an initial value $B^{(0)} \in \mathbb{R}^{N\times N}$, a proximal algorithm solves the optimization problem (2.1) by iteratively calculating the proximal operator of $\Omega = \Omega_{\lambda,\rho}$ for a descent direction of the differentiable loss function ℓ. We use an accelerated version of the algorithm (Beck and Teboulle, 2009), which for each $k = 2, \ldots,$ until convergence, performs the updates

$$\begin{aligned} W^{(k)} &= B^{(k-1)} + \frac{k-1}{k+2}\left(B^{(k-1)} - B^{(k-2)}\right) && (3.1)\\ B^{(k)} &= \text{prox}_{t^{(k)}\Omega}\left\{W^{(k)} - t^{(k)}\nabla\ell\left(W^{(k)}\right)\right\} \\ &= \underset{B\in\mathcal{B}}{\arg\min}\left\{\frac{1}{2}\left\|B - \left(W^{(k)} - t^{(k)}\nabla\ell(W^{(k)})\right)\right\|_2^2 + t^{(k)}\Omega(B)\right\}, && (3.2)\end{aligned}$$

where $\nabla\ell(W) \in \mathbb{R}^{N\times N}$ is the gradient of the loss function ℓ at W and $t^{(k)}, k = 1, 2, \ldots$ is a sequence of positive values. If $\nabla\ell$ is Lipschitz continuous, with L its Lipschitz constant, the sequence of values $\ell\left(B^{(k)}\right) + \Omega\left(B^{(k)}\right)$ converges to the optimal value at rate $O(1/k^2)$ if $t^{(k)} \in [0, 1/L)$. The value of $t^{(k)}$ can be chosen using a backtracking search (Beck and Teboulle, 2009), which decreases this value until the condition

$$\ell\left(B^{(k)}\right) \le \ell\left(W^{(k)}\right) + \left\langle\nabla\ell\left(W^{(k)}\right), B^{(k)} - W^{(k)}\right\rangle + \frac{1}{2t^{(k)}}\|B^{(k)} - W^{(k)}\|_2^2 \quad (3.3)$$

is satisfied. This procedure ensures that step sizes $\{t^{(k)}\}$ become smaller as the algorithm progresses, until $t^{(k)} < 1/L$. In practice, L might be large, which can make the algorithm slow to converge. It has been observed in other sparse high-dimensional problems that search strategies for $t^{(k)}$ which allow for $t^{(k)} > 1/L$ when appropriate can actually speed up convergence (Scheinberg et al., 2014; Hastie et al., 2015). We use a strategy of this type, allowing $t^{(k)}$ to increase by a factor of $\alpha \ge 1$ if the relative improvement in the loss function on iteration k becomes small. We observed that this strategy significantly reduces the number of iterations until convergence. The entire procedure is summarized in Algorithm 1.

Algorithm 1 Proximal algorithm for fitting graph classifier

Input: Training sample $\{(A^{(1)}, Y_1), \ldots, (A^{(n)}, Y_n)\}$; regularization parameters λ, ρ; step size constants $\alpha \ge 1, \delta \in (0, 1), \eta > 0$; tolerance $\epsilon^{\text{PROX}} > 0$.
Initialize: Starting values $B^{(0)}$, $t^{(0)}$.
Iterate: for $k = 1, 2, \ldots$ until $\epsilon^{(k)} < \epsilon^{\text{PROX}}$)

1. Compute $W^{(k)}$ according to (3.1).
2. Compute $B^{(k)}$ by solving the proximal operator (3.2).
3. If condition (3.3) does not hold, decrease step size $t^{(k)} \leftarrow \delta t^{(k)}$ and return to 2.
4. Calculate loss improvement $\epsilon^{(k)} = \left(\ell\left(B^{(k-1)}\right) + \Omega\left(B^{(k-1)}\right)\right) - \left(\ell\left(B^{(k)}\right) + \Omega\left(B^{(k)}\right)\right)$.
5. If $|\epsilon^{(k)} - \epsilon^{(k-1)}|/\epsilon^{(k)} < \eta$, increase step size $t^{(k+1)} = \alpha t^{(k)}$, otherwise set $t^{(k+1)} = t^{(k)}$.

Output: $\hat{B} = B^{(k)}$.

The logistic loss function of (2.4) has an extra parameter b. Rather than including it as an unpenalized coefficient for a constant covariate, we use block coordinate descent and solve for b separately. This is convenient because the threshold b and the matrix of coefficients B may not be on the same scale. Thus, b

can be updated by solving $b^{(k+1)} = \arg\min_{b\in\mathbb{R}} \ell\left(B^{(k)}, b\right)$, which is easy to compute via Netwon's method.

The proximal algorithm requires solving the proximal operator (3.2), which has no closed form solution for the penalty (2.2) under the symmetry constraint. Strategies based on smoothing this penalty have been proposed (Yuan et al., 2011; Chen et al., 2012). However, to allow for variable selection which results from non-differentiability of the penalty, we aim to solve the proximal operator directly using ADMM (see Boyd et al. (2011) for a review). Note that if the symmetric constraint is relaxed as in Remark 1, the proximal operator has a closed form solution (see Remark 4).

The ADMM works by introducing additional constraints and performing coordinate descent in the corresponding augmented Lagrangian function. Setting $Z = W^{(k)} - t^{(k)}\nabla\ell(W^{(k)})$ and $t = t^{(k)}$, and introducing the variables $Q, R \in \mathbb{R}^{N\times N}$, we can formulate (3.2) as a convex optimization problem

$$\begin{aligned} \min_{\tilde{B},Q,R} \quad & \frac{1}{2}\|\tilde{B} - Z\|_2^2 + t\lambda\left(\sum_{i=1}^{N}\|Q_{(i)}\|_2 + \rho\|R\|_1\right) \\ \text{subject to} \quad & \tilde{B} = Q, \quad \tilde{B} = R, \quad \tilde{B} = \tilde{B}^T, \quad \operatorname{diag}(\tilde{B}) = 0. \end{aligned} \tag{3.4}$$

The ADMM algorithm introduces the multipliers $U, V \in \mathbb{R}^{N\times N}$ and a penalty parameter $\mu > 0$ to perform gradient descent on the Lagrangian of (3.4), given by

$$\begin{aligned} \mathcal{L}_\mu(\tilde{B}, Q, R, U, V) \quad = \quad & \frac{1}{2}\left\|\tilde{B} - Z\right\|_2^2 + t\lambda\left(\sum_{i=1}^{N}\|Q_{(i)}\|_2 + \rho\|R\|_1\right) + \left\langle U, \tilde{B} - Q\right\rangle + \left\langle V, \tilde{B} - R\right\rangle \\ & + \frac{\mu}{2}\left(\left\|\tilde{B} - Q\right\|_2^2 + \left\|\tilde{B} - R\right\|_2^2 + \left\|\tilde{B} - R\right\|_2^2 + \left\|\tilde{B} - \tilde{B}^T\right\|_2^2\right). \end{aligned} \tag{3.5}$$

The value μ controls the gap between dual and primal feasibility. In practice, we observed that setting $\mu = 0.1$ gives a good balance between primal and dual feasibility, although other self-tuning methods are available (Parikh and Boyd, 2013). This function is optimized by coordinate descent, with each variable updated to minimize the value of $\mathcal{L}_\mu$ while all the other variables are fixed. This update has a closed form shown in Algorithm 2. These steps are performed until the algorithm converges, to within tolerance $\epsilon^{\text{ADMM}} > 0$ Note that ADMM will be performed in each iteration of the algorithm to solve (2.4) and thus tolerance ϵ^{ADMM} can be decreased as the algorithm progresses. On the other hand, performing only one iteration of algorithm (2) gives a similar algorithm to the one of Chen et al. (2012).

Remark 3. The ADMM makes it very easy to incorporate additional penalties. If Ψ is an new penalty, we can rewrite (3.4) by introducing an additional parameter $\tilde{Q}$ so it becomes

$$\begin{aligned} \min_{\tilde{B},Q,\tilde{Q},R} \quad & \frac{1}{2}\|\tilde{B} - Z^{(k)}\|_2^2 + t\lambda\left(\sum_{i=1}^{N}\|Q_{(i)}\|_2 + \rho\|R\|_1\right) + t\Psi(\tilde{Q}) \\ \text{subject to} \quad & \tilde{B} = Q, \quad \tilde{B} = \tilde{Q}, \quad \tilde{B} = R, \quad \tilde{B} = \tilde{B}^T, \quad \operatorname{diag}(\tilde{B}) = 0. \end{aligned}$$

We can obtain the Lagrangian formulation (3.5) in a similar manner, and include new parameters in the ADMM updates, which can be performed efficiently as long as the proximal operator of Ψ has a closed form solution. This is in fact the case for some other penalties of interest, such as the GraphNet penalty (Grosenick et al., 2013; Watanabe et al., 2014), which can incorporate spatial location information.

Remark 4. The alternative formulation for the graph penalty given in Remark 1 corresponds to standard sparse group lasso (Friedman et al., 2010a). In particular, we can still employ the proximal algorithms (3.1)

Algorithm 2 Proximal operator by ADMM

Input: Z, ϵ^{ADMM}, μ.
Initialize: $\tilde{B}^{(0)} = Z$, $R^{(0)} = Z$, $Q^{(0)} = Z$, $U^{(0)} = 0_{N\times N}$, $V^{(0)} = 0_{N\times N}$.
Iterate: for $l = 1, 2, \ldots$ until convergence ($\epsilon^{(l)}_{\text{ADMM-p}} < \epsilon^{\text{ADMM}}$ and $\epsilon^{(l)}_{\text{ADMM-d}} < \epsilon^{\text{ADMM}}$)

1. Perform coordinate gradient descent on (3.5) by computing

$$\begin{aligned}
\tilde{B}^{(l)} &= \frac{1}{1+2\mu}\left(Z + \frac{1}{2}\mu\left(Q^{(l-1)} + Q^{(l-1)^T}\right) + \mu R^{(l-1)} - U^{(l-1)} - V^{(l-1)}\right) \\
Q^{(l)}_{(i)} &= \left(1 - \frac{t\lambda}{\mu\left\|\tilde{B}^{(l)}_{(i)} + \frac{1}{\mu}U^{(l-1)}_{(i)}\right\|_2}\right)_+ \left(\tilde{B}^{(l)}_{(i)} + \frac{1}{\mu}U^{(l-1)}_{(i)}\right), \qquad i = 1, \ldots, N \\
R^{(l)}_{ij} &= \left(1 - \frac{t\lambda\rho}{\mu\left|\tilde{B}^{(l)}_{ij} + \frac{1}{\mu}V^{(l-1)}_{ij}\right|}\right)_+ \left(\tilde{B}^{(l)}_{ij} + \frac{1}{\mu}V^{(l-1)}_{ij}\right), \quad i = 1, \ldots, N,\ j = 1, \ldots, N, \\
U^{(l)} &= U^{(l-1)} + \mu\left(\tilde{B}^{(l)} - \frac{1}{2}\left(Q^{(l)} + Q^{(l)^T}\right)\right) \\
V^{(l)} &= V^{(l-1)} + \mu\left(\tilde{B}^{(l)} - R^{(l)}\right)
\end{aligned}$$

2. Update primal and dual residuals $\epsilon^{(l)}_{\text{ADMM-p}}$ and $\epsilon^{(l)}_{\text{ADMM-d}}$

$$\begin{aligned}
\epsilon^{(l)}_{\text{ADMM-p}} &= \mu\left(\|Q^{(l)} - Q^{(l-1)}\|_\infty + \|R^{(l)} - R^{(l-1)}\|_\infty\right), \\
\epsilon^{(l)}_{\text{ADMM-d}} &= \mu\left(\|\tilde{B}^{(l)} - Q^{(l)}\|_2 + \|\tilde{B}^{(l)} - R^{(l)}\|_2\right).
\end{aligned}$$

Output: $\tilde{B} = \tilde{B}^{(l)}$.

and (3.2), but instead optimize over the set $\tilde{\mathcal{B}}$. Without the symmetric constraint on B, the overlap in the group lasso penalty dissapears, and this vastly simplifies the problem. Using Theorem 1 of Yuan et al. (2011), the update for $B^{(k)}$ has a closed form solution given by

$$Y^{(k)} = B^{(k-1)} + \frac{k-2}{k}\left(B^{(k-1)} - B^{(k-2)}\right) \tag{3.6}$$

$$Z_{ij}^{(k)} = \left(1 - \frac{\lambda\rho}{\left\|Y_{ij}^{(k)} - t_k \nabla_{ij}\ell(Y^{(k)})\right\|_2}\right)_+ \left(Y_{ij}^{(k)} - t_k \nabla_{ij}\ell(Y^{(k)})\right) \tag{3.7}$$

$$B_{(i)}^{(k)} = \left(1 - \frac{\lambda}{\left\|Z_{(i)}^{(k)}\right\|_2}\right)_+ \left(Z_{(i)}^{(k)}\right), \qquad i = 1, \ldots, N. \tag{3.8}$$

4 Theory

In this section, we show that the solution of the penalized problem (2.2) can recover the correct subgraph corresponding to the set of non-zero coefficients, and give its rates of convergence. These results are based on the framework of Lee et al. (2015) for establishing model selection consistency of regularized M-estimators under geometric decomposability (see Appendix for details).

Let $B^\star \subset \mathbb{R}^{N\times N}$ be the unknown parameter we seek to estimate, and we assume there is a set of active nodes $\mathcal{G} \subset \{1, \ldots, N\}$ with $|\mathcal{G}| = G$, so that $B^\star_{ij} = 0$ if $i \in \mathcal{G}^C$ or $j \in \mathcal{G}^C$. We allow some edge weights inside the subgraph defined by $\mathcal{G}$ to be zero, but we focus on whether the set $\mathcal{G}$ is correctly estimated by the set $\hat{\mathcal{G}}$ of active nodes in $\hat{B}$. Denote by $\mathcal{M} \subseteq \mathbb{R}^{N\times N}$ the set of matrices where the only non-zero coefficients appear in the active subgraph, that is,

$$\mathcal{M} = \left\{B \in \mathbb{R}^{N\times N} \,\middle|\, B_{ij} = 0 \text{ for all } i \in \mathcal{G}^C \text{ or } j \in \mathcal{G}^C, B = B^T\right\} \tag{4.1}$$

There are two main assumptions on the loss function ℓ required for consistent model selection in high-dimensional models (Lee et al., 2015). The first assumption is on the convexity of the loss function around $B^\star$, while the second assumption bounds the size of the entries in the loss Hessian between the variables in the active subgraph and the rest. Let the loss Hessian $\nabla^2\ell(B^\star) \in \mathbb{R}^{N\times N} \otimes \mathbb{R}^{N\times N}$ be defined by

$$\nabla^2_{(i,j),(k,l)}\ell(B) = \frac{\partial^2\ell(B)}{\partial B_{ij}\partial B_{kl}},$$

and define the matrix $H_{(i,j),\mathcal{G}} \in \mathbb{R}^{G\times G}$ with $(i,j) \in (\mathcal{G}\times\mathcal{G})^C$ such that

$$\left(H_{(i,j),\mathcal{G}}\right)_{k,l} = \mathrm{Tr}\left(\left(\nabla^2_{(i,j),(\mathcal{G},\mathcal{G})}\ell(B^\star)\right)\Lambda_{(k,l),(\cdot,\cdot)}\right), \qquad 1 \le k, l \le G, , \tag{4.2}$$

where $\Lambda \in \mathbb{R}^{G\times G} \otimes \mathbb{R}^{G\times G}$ is a tensor such that $\mathrm{Mat}(\Lambda)$ is a pseudoinverse of $\mathrm{Mat}\left(\nabla^2_{(\mathcal{G},\mathcal{G}),(\mathcal{G},\mathcal{G})}\ell(B^\star)\right)$, and Mat is the operation that unfolds the entries of a tensor Λ into a $G^2 \times G^2$ matrix. The matrix $H_{(i,j),\mathcal{G}}$ measures how well the variable corresponding to the edge (i,j) can be represented by the variables in the active subgraph.

Assumption 1 (Restricted Strong Convexity). There exists a set $C \subset \mathbb{R}^{N\times N}$ with $B^\star \in C$, and constants

$m > 0, \tilde{L} < \infty$ such that

$$\begin{aligned} \sum_{i,j} \Delta_{i,j} \mathrm{Tr}\left(\left(\nabla^2_{(i,j),(\cdot,\cdot)} \ell(B)\right) \Delta\right) &\geq m\|\Delta\|_2^2, \quad \text{for all } B \in C \cap \mathcal{M}, \Delta \in C \cap \mathcal{M} \\ \|\nabla^2 \ell(B) - \nabla^2 \ell(B^\star)\|_2 &\leq \tilde{L}\|B - B^\star\|_2, \quad \text{for all } B \in C. \end{aligned}$$

Assumption 2 (Irrepresentability). There exists a constant $0 < \tau < 1$ such that

$$\max_{i \in \mathcal{G}^C} \left\| \left(\sum_{k=1}^{G} \|(H_{(i,j),\mathcal{G}})_{k\cdot}\|_2 \right)_{j=1}^{N} \right\|_2 = 1 - \tau < 0.$$

This version of the irrepresentability condition corresponds to the one usually employed in group lasso penalties (Bach, 2008), but as we will see later, due to overlaps in the groups it further requires a lower bound on ρ to work for model selection.

The first two assumptions are stated directly as a function of the loss for a fixed design case, but they can be substituted with bounds in probability for the case of random designs. In order to obtain rates of convergence, we do require a distributional assumption on the first derivative of the loss. This assumption can be substituted with any bound on $\max_i \|\nabla \ell(B^\star)_{(i)}\|_2$ (see Lemma 3 in the Appendix).

Assumption 3 (Sub-Gaussian score function). Each pair in the sample (A, Y) is independent and comes from a distribution such that the entries of the matrix $\nabla \tilde{\ell}(Y, A; B^\star)$ are subgaussian. That is, for all $t > 0$ there is a constant $\sigma^2 > 0$ such that

$$\max_{i,j} \mathbb{P}\left(\|\nabla_{ij} \tilde{\ell}\,(Y, A; B^\star)\,\|_\infty > t\right) \leq 2 \exp\left(-\frac{t^2}{\sigma^2}\right) .$$

With these assumptions, we establish consistency and correct model selection. The proof is given in the Appendix A.

Proposition 1. Suppose Assumptions 1 and 3 hold. Let $\lambda = c_1 \sqrt{\frac{N \log N}{n}}$ for some constant $c_1 > 0$. Then

(a) The optimal solution of (2.4) is unique and satisfies

$$\|\hat{B} - B^\star\|_2 = O_P\left(\left(\sqrt{N} + \rho N\right) \sqrt{\frac{N \log N}{n}}\right). \tag{4.3}$$

(b) Suppose Assumption 2 also holds. If $n > c_2 \left(\sqrt{G} + \rho G\right) N \log N$ for a constant $c_2 > 0$, and

$$\rho > \frac{1}{\tau} - \frac{1}{\sqrt{G}} , \tag{4.4}$$

then

$$\|\hat{B} - B^\star\|_2 = O_P\left(\rho G \sqrt{\frac{N \log N}{n}}\right), \tag{4.5}$$

$$\mathbb{P}\left(\hat{\mathcal{G}} \subseteq \mathcal{G}\right) = 1 - \exp\left(-c_3 N \log N\right), \tag{4.6}$$

for a constant $c_3 > 0$.

The part of the penalty associated with ρ causes the solution to be sparse. Due to the overlap in the groups, a small value of ρ will not result in zeros in the solution of the problem (2.4). The lower bound on ρ in (4.4) ensures that the irrepresentability condition of Lee et al. (2015) holds (see Lemma 2 in the Appendix). On the other hand, the Frobenius error rate (4.3) is proportional to ρ, and setting ρ equal to zero reduces this bound to $O_P(\sqrt{GN \log N/n})$. This is not uncommon in penalized estimators, where correct support recovery and smallest error rates occur in different regimes.

The result (4.6) ensures that, with high probability, all edges estimated to have non-zero weights are contained in the active subgraph. To ensure that all active nodes are recovered, at least one edge corresponding to each active node needs to have a non-zero weight. A result similar to (4.6) can be obtained to guarantee recovery of all active nodes under a stronger assumption that the magnitude of the non-zero entries of $B^\star$ is bounded below by $|B^\star_{ij}| > c_4 \left(\sqrt{G} + \rho G\right) \lambda$ for a constant c_4.

5 Classification of schizophrenia data

We analyze the performance of the classifier on two different brain fMRI datasets, each containing schizophrenic patients and controls. The first dataset comes from the Center for Biomedical Research Excellence (COBRE). The second dataset, which we refer to as UMich data, is from the lab of Professor Stephan F. Taylor in the Department of Psychiatry at the University of Michigan.

Subjects and imaging – COBRE

Raw anatomic and functional scans from 146 subjects (72 psychosis patients and 74 healthy control subjects) were downloaded from a public database (`http://fcon_1000.projects.nitrc.org/indi/retro/cobre.html`). Four subjects coded as ambidextrous (2 patients, 2 controls) were excluded to yield 70 psychosis patients and 72 controls for analysis. To enter the COBRE dataset, subjects had a diagnosis of either schizophrenia or schizoaffective disorder and were without histories of neurological disorder, mental retardation, severe head trauma with more than 5 minutes loss of consciousness and substance abuse/dependence within the last 12 months.

In the primary sample, two schizophrenic (SCZ) subjects and one healthy control (HC) subject had insufficient voxels in the cerebrospinal fluid (CSF) segmentation on the CSF, and they were dropped from additional analyses. Two additional SCZ subjects were excluded for scrub ratios (see discussion of *scrubbing routine* in fMRI Data Analysis) greater than 0.6, leaving 38 SCZ subjects and 42 HC subjects for the analysis. In the replication sample, 15 psychosis patients and two control subjects were excluded for scrub ratios greater than 0.6; one patient was excluded with incomplete data, leaving 54 SCZ and 69 HC subjects for analysis (see Table 5.2).

A full description of the imaging parameters for the COBRE dataset is available online at the link provided above and in several related papers, see Calhoun et al. (2011); Hanlon et al. (2011); Mayer et al. (2013); Stephen et al. (2013).

Subjects and imaging – UMich data

Subjects were selected from experiments conducted by Professor Stephan F. Taylor at the University of Michigan between 2004 and 2011 for task-based fMRI studies, which included resting state scans. Forty-two

stable outpatients were selected with DSM-IV schizophrenia or schizoaffective disorder (SCZ) (Association et al., 1994). Forty-three healthy comparison (HC) subjects, without a lifetime history of Axis I psychiatric disorders (First et al., 1995), were selected to approximate the age range, gender distribution and family education level of the patients. Prior to initial data collection, all subjects gave written, informed consent to participate in the protocol approved by the University of Michigan institutional review board (IRBMED).

MRI scanning occurred on a GE 3T Signa scanner (LX [8.3] release, General Electric Healthcare, Buckinghamshire, United Kingdom). Functional images were acquired with a T2*-weighted, reverse spiral acquisition sequence (gradient recalled echo, TE=30 msec, FA=90 degrees, field of view=22 cm, 40 slices, 3.0mm thick/0mm skip, equivalent to 64 x 64 voxel grid – yielding isotropic voxels 3 mm on edge). Because the data were acquired across different experiments, acquisition parameters differed slightly in the aggregate sample: 240 volumes @ TR=1500 msec (11 SCZ, 10 HC), 180 volumes @ TR=2000 msec (17 SCZ, 16 HC) and 240 volumes at 2000 msec (14 SCZ, 17 HC). Acquisitions were acquired in the resting state with eyes open and fixated on a large 'plus' sign projected on a monitor.

fMRI Data Analysis

Data processing began with standard pre-processing steps. All scans were slice-time corrected and realigned to the 10th image acquired during a scanning session (Jenkinson et al., 2002). Subsequent processing was performed with the Statistical Parametric Mapping SPM8 package (Wellcome Institute of Cognitive Neurology, London). Anatomic normalization was done with the VBM8 toolbox in SPM8, using the high resolution structural scans obtained for both datasets. Normalizing warps were applied to the co-registered, functional volumes, which were re-sliced and smoothed with an 8 mm isotropic Gaussian smoothing kernel. To assess and manage movement, we calculated the frame-wise displacement (FD) (Power et al., 2012), for all 6 parameters of rotation and translation. We used a *scrubbing routine* to censor any frame with FD > 0.5 mm from the regression analysis described below, yielding a scrub ratio for each subject. Three-compartment segmentation of the high-resolution structural image from the VBM8 normalization was applied to the functional time series to extract cerebral spinal volume (CSF) and white matter (WM) compartments, which were then subjected to a principal component analysis to identify the top 5 components in each (Behzadi et al., 2007), which should correspond to heart rate and respiratory effects on global signal (Chai et al., 2012). Multiple regressions were applied to the time series to remove the following nuisance effects: Linear trend, 6 motion parameters, their temporal derivatives, the quadratics of these 12 parameters, 5 components from the PCA of CSF, 5 components of PCA of WM, followed by band pass filtering from 0.01 – 0.1 Hz, and then motion scrubbing. For each 4D data set, time courses were then extracted from 10 mm diameter spheres based on the 264 sets of coordinates identified by Power and colleagues (Power et al., 2011). From these time series, a cross-correlation matrix of Pearson r-values was obtained and Fisher's R-to-Z transformation was applied for each of the 264 nodes with every other node (for COBRE dataset, node 75 is missing). Finally, edge weights were assigned to be ranks of these score, with edge scores ranked separately for each subject. Ranks have been used previously in brain connectomic studies to reduce the effect of potential outliers (Yan et al., 2013); we observed that while ranks does not increase the classification accuracy significantly, they tend to produce sparser solutions with a similar accuracy to Pearson correlations.

First, we evaluate the performance of the classifier using 10-fold cross-validation to choose tuning parameters and to estimate the test error. The classifier is trained for a range of values of λ and ρ, with $\lambda \in \left\{10^{-7}, 10^{-6.5}, \ldots, 10^{-2}\right\}$ and $\rho \in \left\{10^{-3}, 10^{-2.5}, \ldots, 10^{2}\right\}$. The value of γ in (2.4) is set to 10^{-5}; we observed that setting γ to a small value does not affect the accuracy or sparsity of the solution, but it speeds

	System	Number of nodes
1	Sensory/somatomotor Hand	30
2	Sensory/somatomotor Mouth	5
3	Cingulo-opercular Task Control	14
4	Auditory	13
5	Default mode	58
6	Memory retrieval	5
7	Visual	31
8	Fronto-parietal Task Control	25
9	Salience	18
10	Subcortical	13
11	Ventral attention	9
12	Dorsal attention	11
13	Cerebellar	4
-1	Uncertain	28

Table 5.1: Functional systems for the brain parcellation of Power et al. (2011).

Dataset	# nodes	Status	# patients	Male/female	Average age (s.d.)
COBRE	263	Schizophrenic	54	48/6	35.4 (13.1)
		Control	70	48/22	35.1 (11.5)
UMich	264	Schizophrenic	39	29/10	40.7 (11.5)
		Control	40	28/12	36.8 (12.3)

Table 5.2: Summary of the two datasets.

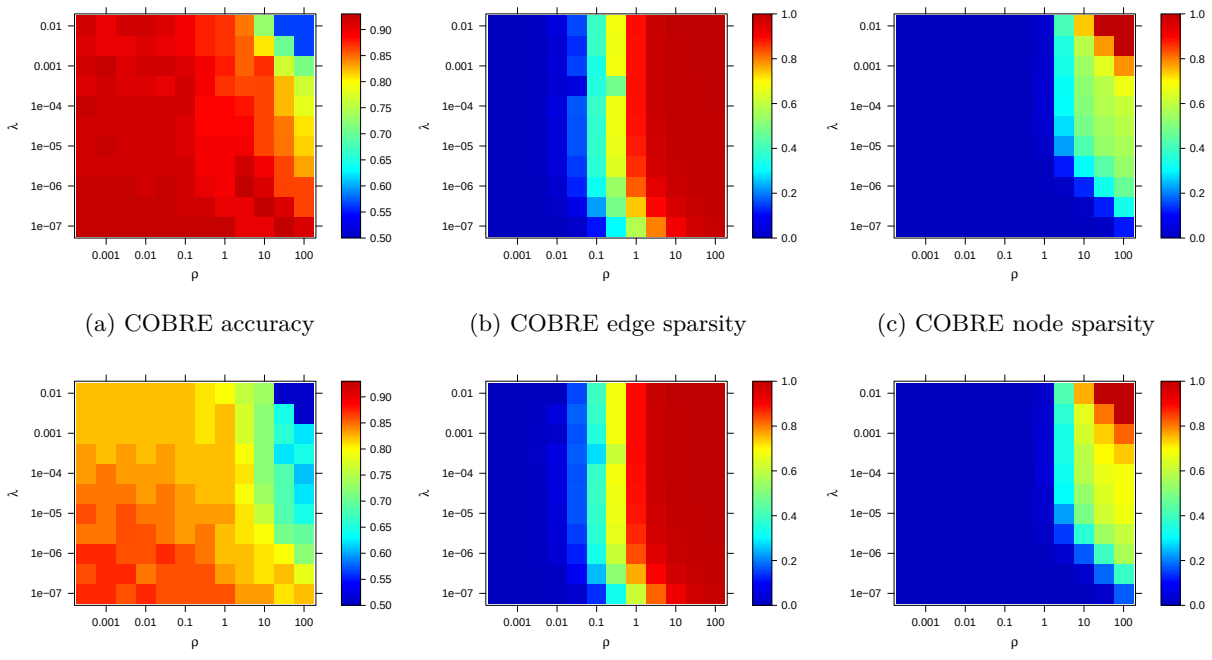

(a) COBRE accuracy (b) COBRE edge sparsity (c) COBRE node sparsity

(d) UMich accuracy (e) UMich edge sparsity (f) UMich node sparsity

Figure 5.1: Cross-validated results for the two data sets. Classifcation accuracy (left), fraction of zero edge coefficients (middle), and fraction of inactive nodes (right).

Method	Classification accuracy % (standard deviation)	
	COBRE	UMich
Our method	92.7 (2.6)	85.9 (3.6)
Lasso	80.1 (5.6)	60.9 (5.6)
Elastic net	89.5 (1.8)	82.6 (4.7)
Ridge penalty	91 (2.6)	80.9 (3.5)

Table 5.3: Cross-validated accuracy (average and standard deviation over folds) o for different methods.

up convergence. Figure 5.1 shows the average cross-validated accuracy, sparsity (fraction of zero coefficients) and node sparsity (fraction of inactive nodes), as a heatmap over tuning parameter values. for each tuning parameter on all the cross-validation subsets. As expected, accuracy generally decreases as the solution becomes sparser, which is not uncommon in high-dimensional settings (Hastie et al., 2015). However, we can still achieve excellent accuracy with a substantially reduced set of features. We also observe that λ has little influence on sparsity, which is primarily controlled by ρ. Moreover, as Proposition 1 suggests, values of $\rho < 1$ do not result in node selection.

For comparison, we include several general linear classifiers with non-structured regularization penalties: the lasso, the elastic net, and ridge (Friedman et al., 2010b), fitted using the `R` package `glmnet`. For the elastic net, we fix the second tuning parameter $\alpha = 0.2$. We perform 10-fold cross-validation to estimate the accuracy of each classifier, with a nested cross-validation to pick the best parameters for each fold. Table 5.3 shows the average cross-validated accuracy for each method, together with its standard error. Our structured penalty achieves higher accuracy than all the non-structured classifiers. All methods perform better on the COBRE dataset than on the UMich dataset, which can be partially explained by the different sample sizes and possibly noise levels. Comparing with other methods in the literature, Watanabe et al. (2014) evaluated the performance of their classifiers on a different parcellation of the COBRE data (since their methods are based on the assumption of equally spaced nodes, they cannot be directly applied on the parcellation we used). Although this makes the results not directly comparable, it is the same dataset, and our method, besides being more general, substantially outperforms the methods of Watanabe et al. (2014), whose reported accuracy was between 71.9% and 73.5%.

Besides prediction, a meaningful goal is to select a subset of important variables. Since best classification performance here is achieved with all variables included, we use the so called "one-standard-error rule" (Hastie et al., 2015), who recommend selecting the most parsimonious classifier with cross-validation accuracy at most one standard error away from the best cross-validation accuracy. We measure parsimony by the number of inactive nodes. Using this rule, we obtain the average cross-validated classification accuracy of 89.4% (COBRE) and 82.1% (UMich), which is not far from the best classification accuracy shown in Table 5.3. Figure 5.2 shows the solutions for each dataset obtained by this rule. Nodes are ordered by brain systems of Table 5.1. The fitted solution for COBRE has 1098 non-zero coefficients (1.56%) and 217 active nodes (82.5%), while the UMich solution has 23496 non-zero entries (45.8%), but all nodes are active. Note that since the networks are undirected, each coefficient is repeated twice, so the total number of different coefficients is half than that.

Ultimately, assessing significance of the selected variables is necessary, which is in general a difficult task in high-dimensional settings and an active area of research (see for example Meinshausen and Bühlmann (2010); Van de Geer et al. (2014); Lockhart et al. (2014)). In brain connectomics, it is particularly challenging to identify significant variables because of small sample sizes (Button et al., 2013). Here we employ *stability selection* (Meinshausen and Bühlmann, 2010) which can be shown to control a type of false discovery rate.

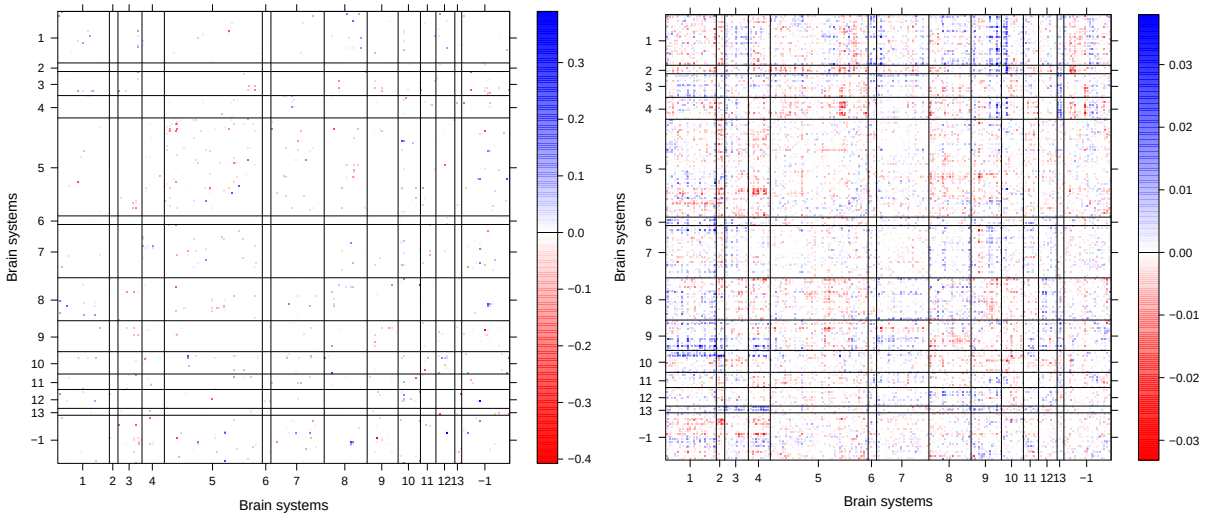

(a) COBRE solution

(b) UMich solution

Figure 5.2: Fitted coefficients for COBRE and UMich datasets, with tuning parameters selected by the "one standard-error".

The method works by estimating the probability of selecting a variable by splitting the data. Some versions of this method have been theoretically studied, and upper bounds on the expected number of variables with a low selection probability that are included in the final solution have been derived under mild conditions (Meinshausen and Bühlmann, 2010; Shah and Samworth, 2013); under stronger conditions this can be related to the number of false discoveries. In particular, we use simultaneous selection (Shah and Samworth, 2013), which we empirically observed does better in controlling small selection probabilities. The number of random splits is set to $m = 50$. As seen in Figure 5.1, the value of λ has little effect on the sparsity patterns, so we fix the value of $\lambda = 10^{-4}$ but vary ρ so that $\rho \in \{10^0, 10^{0.2}, \ldots, 10^2\}$.

We estimate simultaneous selection probabilities for each variable on both datasets, shown in the top row of Figure 5.3. These probabilities may be easier to interpret for a group of edges. The middle row of Figure 5.3 shows the maximum probability of selection among the edges attached to each node, which can be interpreted as a lower bound on the probability of activating that node. Similarly, the bottom row shows the maximum probability of selection for each cell (a cell is a group of edges connecting two functional systems listed in Table 5.1, or a group of edges within a single region, so we have a total of $\binom{14}{2} + 14 = 105$ distinct cells. The cells that contain edges with an estimated selection probability of at least 0.9 are highlighted with red in the bottom row of Figure 5.3. Shah and Samworth (2013) provide an upper bound on the expected number of falsely selected variables with this method, which under further conditions can be related to the expected number of false discoveries (Meinshausen and Bühlmann, 2010). Stability selection chooses 22 different edges in 19 cells in the COBRE dataset, and 12 different edges in 10 cells in the UMich dataset. Three of these cells (sensory/somatomotor hand with default mode and fronto-parietal task control systems, and auditory system with uncertain nodes) coincide in both datasets, although all edges selected are different. Theorem 1 and Lemma 1 of Shah and Samworth (2013) indicate that the expected number of falsely selected variables out of these is 6.15 for COBRE and 5.25 for UMich data, which suggests

the majority of selected cells are indeed important. In both datasets, cells in the default mode network (5) and fronto-parietal task control region (8) especially stand out. These systems have been previously linked to schizophrenia (Bunney and Bunney, 2000; Fornito et al., 2012). These results coincide with the findings of Watanabe et al. (2014) on a different parcellation of the same data, which is an encouraging indication of robustness of the network representation to the exact choice of node locations. Some of the variables with the highest estimated selection probabilities appear in the uncertain system (-1), in particular in the cell connecting it with salience system (9), which suggests that alternative parcellations that better characterize these regions may offer a better account of the schizophrenia-related changes identified therein. Additionally, in the UMich data, the sensory/somatomotor system appears to be important. The results from UMich data may be less reliable due to the smaller sample size and substantially lower classification accuracy as reported in Table 5.3.

6 Discussion

We have presented a method for classifying graphs with labeled nodes, motivated by brain connectomics but generally applicable to any setting with such graphs. The distinct feature of our method is that it is graph-aware, aiming for a sparse set of *both* edges and nodes selected, but it is general in the sense that it does not rely on the spatial structure specific to the brain. The method is computationally efficient since the regularization we use is convex.

The results we obtained on the schizophrenia data are in agreement with the consistent findings of previous studies on the default mode network, which is one of the networks implicated across a wide variety of psychiatric disorders (Öngür et al., 2010; Broyd et al., 2009). The findings are also notable for the divergent findings between the two datasets, wherein non-overlapping different networks were implicated in each dataset. These differences may reflect real differences in samples collected in different sites, as significant pathophysiological heterogeneity occurs for all psychiatric diagnoses, or they may simply reflect type 2 errors.

Acknowledgments

E. Levina's research is partially supported by NSF grant DMS-1521551 and ONR grant N000141612910. S. F. Taylor's research is supported by the National Institute of Mental Health (R01MH064148, R21MH086701, R21MH101676), the Boledovich Schizophrenia Research Fund and University of Michigan Clinical Translational Science Award (UL1RR024986).

References

Association, A. P., Association, A. P., et al. (1994). Diagnostic and statistical manual of mental disorders (DSM). *Washington, DC: American psychiatric association*, pages 143–7.

Bach, F., Jenatton, R., Mairal, J., Obozinski, G., et al. (2012). Structured sparsity through convex optimization. *Statistical Science*, 27(4):450–468.

Bach, F. R. (2008). Consistency of the group lasso and multiple kernel learning. *Journal of Machine Learning Research*, 9(Jun):1179–1225.

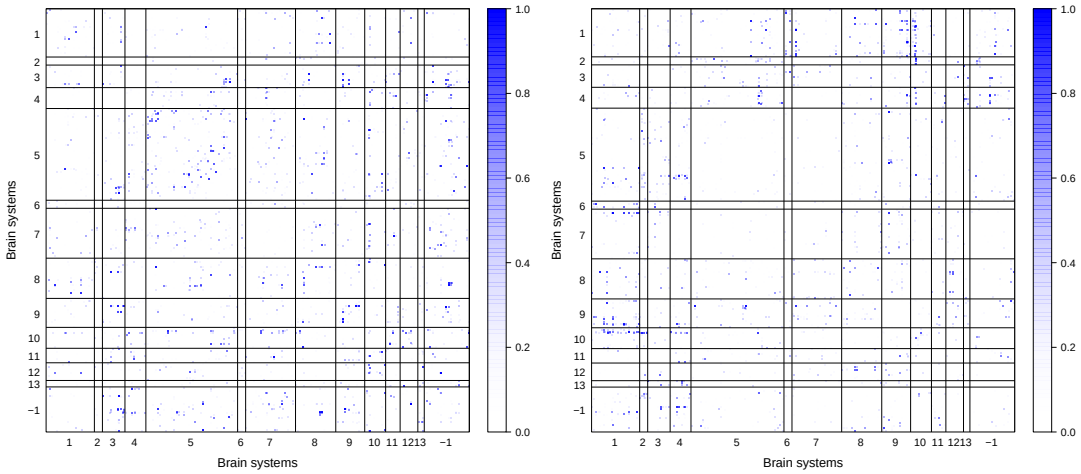

(a) Edge selection probabilities (COBRE)

(b) Edge selection probabilities (UMich)

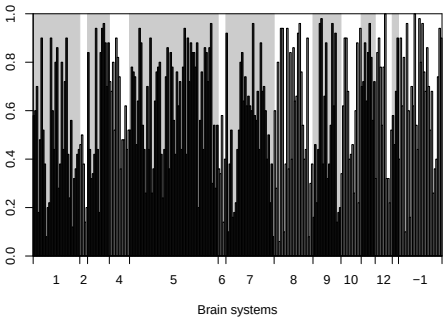

(c) Node selection probabilities (COBRE)

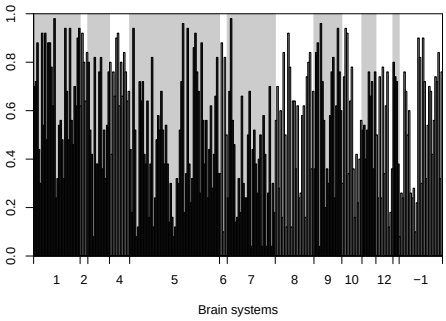

(d) Node selection probabilities (UMich)

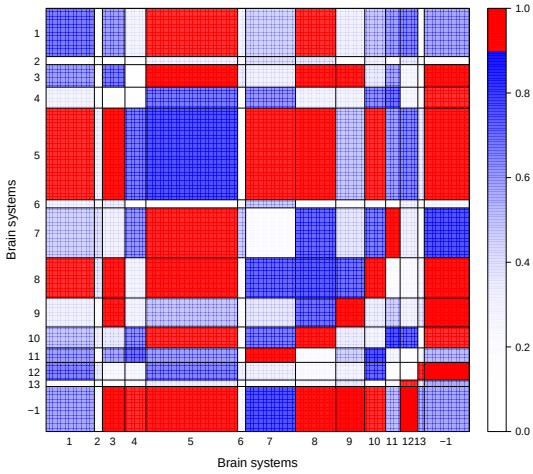

(e) Cell selection probabilities (COBRE)

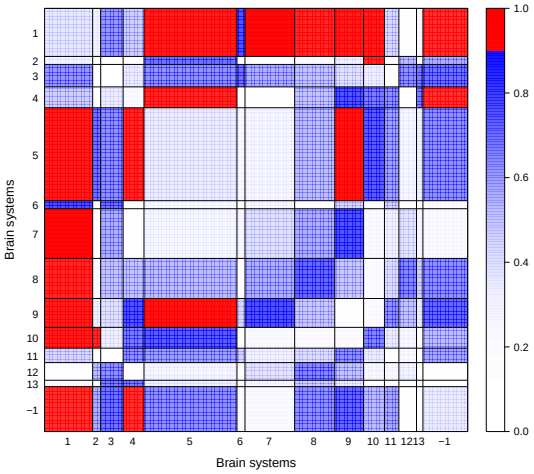

(f) Cell selection probabilities (UMich)

Figure 5.3: Estimated selection probability for each edge (top row), as well as estimated lower bounds for selecting at least one edge of each node (middle row) or each cell (bottom row). The probabilities are calculated using simultatenous stability selection (Shah and Samworth, 2013) with 50 random splits.

Beck, A. and Teboulle, M. (2009). A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202.

Behzadi, Y., Restom, K., Liau, J., and Liu, T. T. (2007). A component based noise correction method (compcor) for bold and perfusion based fMRI. *Neuroimage*, 37(1):90–101.

Bickel, P. J. and Chen, A. (2009). A nonparametric view of network models and newman–girvan and other modularities. *Proceedings of the National Academy of Sciences*, 106(50):21068–21073.

Boyd, S., Parikh, N., Chu, E., Peleato, B., and Eckstein, J. (2011). Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122.

Broyd, S. J., Demanuele, C., Debener, S., Helps, S. K., James, C. J., and Sonuga-Barke, E. J. (2009). Default-mode brain dysfunction in mental disorders: a systematic review. *Neuroscience & biobehavioral reviews*, 33(3):279–296.

Bullmore, E. and Sporns, O. (2009). Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience*, 10(3):186–198.

Bunney, W. E. and Bunney, B. G. (2000). Evidence for a compromised dorsolateral prefrontal cortical parallel circuit in schizophrenia. *Brain Research Reviews*, 31(2):138–146.

Button, K. S., Ioannidis, J. P., Mokrysz, C., Nosek, B. A., Flint, J., Robinson, E. S., and Munafò, M. R. (2013). Power failure: why small sample size undermines the reliability of neuroscience. *Nature Reviews Neuroscience*, 14(5):365–376.

Calhoun, V. D., Sui, J., Kiehl, K., Turner, J., Allen, E., and Pearlson, G. (2011). Exploring the psychosis functional connectome: aberrant intrinsic networks in schizophrenia and bipolar disorder. *Magnetic resonance imaging of disturbed brain connectivity in psychiatric illness*, 35.

Chai, X. J., Castañón, A. N., Öngür, D., and Whitfield-Gabrieli, S. (2012). Anticorrelations in resting state networks without global signal regression. *Neuroimage*, 59(2):1420–1428.

Chen, X., Lin, Q., Kim, S., Carbonell, J. G., and Xing, E. P. (2012). Smoothing proximal gradient method for general structured sparse regression. *The Annals of Applied Statistics*, 6(2):719–752.

Craddock, R. C., Holtzheimer, P. E., Hu, X. P., and Mayberg, H. S. (2009). Disease state prediction from resting state functional connectivity. *Magnetic Resonance in Medicine*, 62(6):1619–1628.

Deshpande, M., Kuramochi, M., Wale, N., and Karypis, G. (2005). Frequent substructure-based approaches for classifying chemical compounds. *IEEE Transactions on Knowledge and Data Engineering*, 17(8):1036–1050.

Duchi, J. and Singer, Y. (2009). Efficient online and batch learning using forward backward splitting. *Journal of Machine Learning Research*, 10(Dec):2899–2934.

Fei, H. and Huan, J. (2010). Boosting with structure information in the functional space: an application to graph classification. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 643–652. ACM.

First, M. B., Spitzer, R. L., Gibbon, M., Williams, J. B., et al. (1995). Structured clinical interview for DSM-IV axis i disorders. *New York: New York State Psychiatric Institute*.

Fornito, A., Zalesky, A., Pantelis, C., and Bullmore, E. T. (2012). Schizophrenia, neuroimaging and connectomics. *NeuroImage*, 62(4):2296–2314.

Friedman, J., Hastie, T., and Tibshirani, R. (2010a). A note on the group lasso and a sparse group lasso. *arXiv preprint arXiv:1001.0736*.

Friedman, J., Hastie, T., and Tibshirani, R. (2010b). Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33(1):1.

Gao, C., Lu, Y., Zhou, H. H., et al. (2015). Rate-optimal graphon estimation. *The Annals of Statistics*, 43(6):2624–2652.

Ginestet, C. E., Balanchandran, P., Rosenberg, S., and Kolaczyk, E. D. (2014). Hypothesis testing for network data in functional neuroimaging. *arXiv preprint arXiv:1407.5525*.

Gonzalez, J., Holder, L. B., and Cook, D. J. (2000). Graph based concept learning. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence and Twelfth Conference on Innovative Applications of Artificial Intelligence. AAAI Press*, page 1072. The MIT Press.

Gotts, S., Saad, Z., Jo, H. J., Wallace, G., Cox, R., and Martin, A. (2013). The perils of global signal regression for group comparisons: a case study of Autism Spectrum Disorders. *Frontiers in Human Neuroscience*, 7:356.

Grosenick, L., Klingenberg, B., Katovich, K., Knutson, B., and Taylor, J. E. (2013). Interpretable whole-brain prediction analysis with graphnet. *NeuroImage*, 72:304–321.

Hanlon, F. M., Houck, J. M., Pyeatt, C. J., Lundy, S. L., Euler, M. J., Weisend, M. P., Thoma, R. J., Bustillo, J. R., Miller, G. A., and Tesche, C. D. (2011). Bilateral hippocampal dysfunction in schizophrenia. *Neuroimage*, 58(4):1158–1168.

Hastie, T., Tibshirani, R., and Wainwright, M. (2015). *Statistical Learning with Sparsity: The Lasso and Generalizations*. CRC Press.

Helma, C., King, R. D., Kramer, S., and Srinivasan, A. (2001). The predictive toxicology challenge 2000–2001. *Bioinformatics*, 17(1):107–108.

Hu, Y. and Allen, G. I. (2015). Local-aggregate modeling for big data via distributed optimization: Applications to neuroimaging. *Biometrics*, 71(4):905–917.

Inokuchi, A., Washio, T., and Motoda, H. (2000). An apriori-based algorithm for mining frequent substructures from graph data. In *Principles of Data Mining and Knowledge Discovery*, pages 13–23. Springer.

Jacob, L., Obozinski, G., and Vert, J.-P. (2009). Group lasso with overlap and graph lasso. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 433–440. ACM.

Jenkinson, M., Bannister, P., Brady, M., and Smith, S. (2002). Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage*, 17(2):825–841.

Kashima, H., Tsuda, K., and Inokuchi, A. (2003). Marginalized kernels between labeled graphs. In *International Conference of Machine Learning*, volume 3, pages 321–328.

Ketkar, N. S., Holder, L. B., and Cook, D. J. (2009). Empirical comparison of graph classification algorithms. In *Computational Intelligence and Data Mining, 2009. CIDM'09. IEEE Symposium on*, pages 259–266. IEEE.

Kudo, T., Maeda, E., and Matsumoto, Y. (2004). An application of boosting to graph classification. In *Advances in Neural Information Processing Systems*, pages 729–736.

Le, C. M., Levina, E., and Vershynin, R. (2015). Sparse random graphs: regularization and concentration of the Laplacian. *arXiv preprint arXiv:1502.03049*.

Lee, J. D., Sun, Y., Taylor, J. E., et al. (2015). On model selection consistency of regularized m-estimators. *Electronic Journal of Statistics*, 9:608–642.

Liu, Y., Liang, M., Zhou, Y., He, Y., Hao, Y., Song, M., Yu, C., Liu, H., Liu, Z., and Jiang, T. (2008). Disrupted small-world networks in schizophrenia. *Brain*, 131(4):945–961.

Lockhart, R., Taylor, J., Tibshirani, R. J., and Tibshirani, R. (2014). A significance test for the lasso. *The Annals of Statistics*, 42(2):413.

Mayer, A. R., Ruhl, D., Merideth, F., Ling, J., Hanlon, F. M., Bustillo, J., and Cañive, J. (2013). Functional imaging of the hemodynamic sensory gating response in schizophrenia. *Human brain mapping*, 34(9):2302–2312.

Meinshausen, N. and Bühlmann, P. (2006). High-dimensional graphs and variable selection with the lasso. *The Annals of Statistics*, 34(3):1436–1462.

Meinshausen, N. and Bühlmann, P. (2010). Stability selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(4):417–473.

Narayan, M., Allen, G. I., and Tomson, S. (2015). Two sample inference for populations of graphical models with applications to functional connectivity. *arXiv preprint arXiv:1502.03853*.

Öngür, D., Lundy, M., Greenhouse, I., Shinn, A. K., Menon, V., Cohen, B. M., and Renshaw, P. F. (2010). Default mode network abnormalities in bipolar disorder and schizophrenia. *Psychiatry Research: Neuroimaging*, 183(1):59–68.

Parikh, N. and Boyd, S. (2013). Proximal algorithms. *Foundations and Trends in Optimization*, 1(3):123–231.

Power, J. D., Barnes, K. A., Snyder, A. Z., Schlaggar, B. L., and Petersen, S. E. (2012). Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. *Neuroimage*, 59(3):2142–2154.

Power, J. D., Cohen, A. L., Nelson, S. M., Wig, G. S., Barnes, K. A., Church, J. A., Vogel, A. C., Laumann, T. O., Miezin, F. M., Schlaggar, B. L., et al. (2011). Functional network organization of the human brain. *Neuron*, 72(4):665–678.

Richiardi, J., Eryilmaz, H., Schwartz, S., Vuilleumier, P., and Van De Ville, D. (2011). Decoding brain states from fMRI connectivity graphs. *NeuroImage*, 56(2):616–626.

Scheinberg, K., Goldfarb, D., and Bai, X. (2014). Fast first-order methods for composite convex optimization with backtracking. *Foundations of Computational Mathematics*, 14(3):389–417.

Shah, R. D. and Samworth, R. J. (2013). Variable selection with error control: another look at stability selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 75(1):55–80.

Smith, S. M., Vidaurre, D., Beckmann, C. F., Glasser, M. F., Jenkinson, M., Miller, K. L., Nichols, T. E., Robinson, E. C., Salimi-Khorshidi, G., Woolrich, M. W., et al. (2013). Functional connectomics from resting-state fMRI. *Trends in Cognitive Sciences*, 17(12):666–682.

Srinivasan, A., Muggleton, S. H., Sternberg, M. J., and King, R. D. (1996). Theories for mutagenicity: A study in first-order and feature-based induction. *Artificial Intelligence*, 85(1):277–299.

Sripada, C., Angstadt, M., Kessler, D., Phan, K. L., Liberzon, I., Evans, G. W., Welsh, R. C., Kim, P., and Swain, J. E. (2014a). Volitional regulation of emotions produces distributed alterations in connectivity between visual, attention control, and default networks. *Neuroimage*, 89:110–121.

Sripada, C., Kessler, D., Fang, Y., Welsh, R. C., Prem Kumar, K., and Angstadt, M. (2014b). Disrupted network architecture of the resting brain in attention-deficit/hyperactivity disorder. *Human brain mapping*, 35(9):4693–4705.

Stephen, J. M., Coffman, B. A., Jung, R. E., Bustillo, J. R., Aine, C., and Calhoun, V. D. (2013). Using joint ICA to link function and structure using MEG and DTI in schizophrenia. *Neuroimage*, 83:418–430.

Supekar, K., Menon, V., Rubin, D., Musen, M., and Greicius, M. D. (2008). Network analysis of intrinsic functional brain connectivity in alzheimer's disease. *PLoS Comput Biol*, 4(6):e1000100.

Tang, M., Athreya, A., Sussman, D. L., Lyzinski, V., and Priebe, C. E. (2014a). A nonparametric two-sample hypothesis testing problem for random dot product graphs. *arXiv preprint arXiv:1409.2344*.

Tang, M., Athreya, A., Sussman, D. L., Lyzinski, V., and Priebe, C. E. (2014b). A semiparametric two-sample hypothesis testing problem for random dot product graphs. *arXiv preprint arXiv:1403.7249*.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society; Series B*, 73(3):267–288.

Van de Geer, S., Bühlmann, P., Ritov, Y., Dezeure, R., et al. (2014). On asymptotically optimal confidence regions and tests for high-dimensional models. *The Annals of Statistics*, 42(3):1166–1202.

Varoquaux, G. and Craddock, R. C. (2013). Learning and comparing functional connectomes across subjects. *NeuroImage*, 80:405–415.

Vogelstein, J. T., Roncal, W. G., Vogelstein, R. J., and Priebe, C. E. (2013). Graph classification using signal-subgraphs: Applications in statistical connectomics. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(7):1539–1551.

Watanabe, T., Kessler, D., Scott, C., Angstadt, M., and Sripada, C. (2014). Disease prediction based on functional connectomes using a scalable and spatially-informed support vector machine. *NeuroImage*, 96:183–202.

Xin, B., Kawahara, Y., Wang, Y., and Gao, W. (2014). Efficient generalized fused lasso and its application to the diagnosis of alzheimer's disease. In *Twenty-Eighth AAAI Conference on Artificial Intelligence*, pages 2163–2169.

Yan, C.-G., Craddock, R. C., Zuo, X.-N., Zang, Y.-F., and Milham, M. P. (2013). Standardizing the intrinsic brain: towards robust measurement of inter-individual variation in 1000 functional connectomes. *Neuroimage*, 80:246–262.

Yuan, L., Liu, J., and Ye, J. (2011). Efficient methods for overlapping group lasso. In *Advances in Neural Information Processing Systems*, pages 352–360.

Yuan, M. and Lin, Y. (2006). Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society; Series B*, 68(1):49–67.

Zalesky, A., Fornito, A., and Bullmore, E. T. (2010). Network-based statistic: identifying differences in brain networks. *NeuroImage*, 53(4):1197–1207.

Zhang, A. Y. and Zhou, H. H. (2015). Minimax rates of community detection in stochastic block models. *arXiv preprint arXiv:1507.05313*.

Zhang, J., Cheng, W., Wang, Z., Zhang, Z., Lu, W., Lu, G., and Feng, J. (2012). Pattern classification of large-scale functional brain networks: identification of informative neuroimaging markers for epilepsy. *PLoS One*, 7(5):e36733.

A Appendix

Here we prove the bounds on Frobenius norm error and probability of support selection in Proposition 1, following the framework of Lee et al. (2015) based on geometrical decomposability. A penalty Ω is geometrically decomposable if it can be written as

$$\Omega(B) = h_A(B) + h_I(B) + h_{E^\perp}(B)$$

for all B, with A, I closed convex sets, E a subspace, and $h_C(B) = \sup\{\langle Y, B\rangle \,|Y \in C\}$ the support function on C.

The proof proceeds in the following steps. Lemma 1 shows that an equivalent form of our penalty (2.2) is geometrically decomposable, allowing us to use the framework of Lee et al. (2015). Lemma 3 shows the Assumption 2 together with a lower bound on ρ imply that the irrepresentability assumption of Lee et al. (2015) holds. Assumption 2 is directly on the entries of the loss Hessian, which simplifies the very general form of the assumption in Lee et al. (2015). Lemma 3 gives a bound on the entries of the loss gradient under the sub-Gaussianity assumption 3. Lemma 4 gives explicit bounds for the compatibility constants that appear on Theorem 1 of Lee et al. (2015). Finally, we combine these results to prove Proposition 1.

Without loss of generality, to simplify notation we assume that $\mathcal{G} = \{1, \ldots, G\}$, that is, the active subgraph is in the first G rows of the matrix.

Lemma 1. *The penalty* (2.2) *can be written as geometrically decomposable.*

Proof of Lemma 1. We use an equivalent formulation of the penalty in which every coefficient is penalized only once. Let $B^{(1)}, B^{(2)} \in \mathbb{R}^{N\times N}$ be matrices such that the upper triangular part of $B^{(2)}$ and the diagonals of $B^{(1)}$ and $B^{(2)}$ are zero. Define

$$\tilde{\Omega}\left(B^{(1)}, B^{(2)}\right) = \sum_{i=1}^{N} \|B^{(1)}_{(i)}\|_2 + \rho\|B^{(1)}\|_1,$$

$$E = \{(B^{(1)}, B^{(2)}) \in \mathbb{R}^{N\times 2N} \ : \ B^{(1)} = {B^{(1)}}^T, B^{(2)}_{ij} = B^{(1)}_{ij}, \text{ for } i < j \text{ and } B^{(2)}_{ij} = 0 \text{ for } i \geq j\}.$$

Denote by R the transformation from $\mathbb{R}^{N\times N}$ to $\mathbb{R}^{N\times 2N}$ that replicates entires appropriately,

$$(RB)_{ij} = \begin{cases} B_{ij} & \text{if } 1 \leq j \leq N \\ B_{i(j-N)} & \text{if } j > N. \end{cases} \tag{A.1}$$

Therefore, for any $B \in \mathbb{R}^{N\times N}$, we can uniquely define $RB = (B^{(1)}, B^{(2)})$ such that $\Omega(B) = \tilde{\Omega}(B^{(1)}, B^{(2)})$. We then show that $\tilde{\Omega}$ is geometrically decomposable. Moreover, for any $(B^{(1)}, B^{(2)}) \in E$ we can define R^{-1}, so the penalties Ω and $\tilde{\Omega}$ on E are equivalent. Define the sets $A, I \subset \mathbb{R}^{N\times 2N}$ such that

$$\begin{aligned} A &= \left\{(B^{(1)}, B^{(2)}) : \max_{i\in\mathcal{G}} \|B^{(1)}_{(i)}\|_2 \leq 1, \max_{i\in\mathcal{G}^C} \|B^{(1)}_{(i)}\|_2 = 0, \ \max|B^{(2)}_{ij}| \leq \rho, B^{(2)}_{ij} = 0, (i,j) \in (\mathcal{G}\times\mathcal{G})^C\right\}, \\ I &= \left\{(B^{(1)}, B^{(2)}) : \max_{i\in\mathcal{G}^C} \|B^{(1)}_{(i)}\|_2 \leq 1, \max_{i\in\mathcal{G}} \|B^{(1)}_{(i)}\|_2 = 0, \ \max|B^{(2)}_{ij}| \leq \rho, B^{(2)}_{ij} = 0, (i,j) \in \mathcal{G}\times\mathcal{G}\right\}. \end{aligned}$$

Letting $\left\langle Y, (B^{(1)}, B^{(2)})\right\rangle = \mathrm{Tr}(Y^{(1)}{B^{(1)}}^T) + \mathrm{Tr}(Y^{(2)}{B^{(2)}}^T)$, combining the arguments of Lee et al. (2015) for lasso and group lasso penalties,

$$\begin{aligned} h_A\left(B^{(1)}, B^{(2)}\right) &= \sum_{i\in\mathcal{G}} \|B^{(1)}_{(i)}\|_2 + \rho \sum_{(i,j)\in\mathcal{G}\times\mathcal{G}} |B^{(2)}_{ij}|, \\ h_I\left(B^{(1)}, B^{(2)}\right) &= \sum_{i\in\mathcal{G}^C} \|B^{(1)}_{(i)}\|_2 + \rho \sum_{(i,j)\in(\mathcal{G}\times\mathcal{G})^C} |B^{(2)}_{ij}|, \\ h_E\left(B^{(1)}, B^{(2)}\right) &= \begin{cases} 0 & \text{if } \left(B^{(1)}, B^{(2)}\right) \in E \\ \infty & \text{otherwise.} \end{cases} \end{aligned}$$

Hence, Ω can be written as a geometrically decomposable penalty

$$\Omega(B) = \tilde{\Omega}\left(B^{(1)}, B^{(2)}\right) = \lambda\left(h_A\left(B^{(1)}, B^{(2)}\right) + h_I\left(B^{(1)}, B^{(2)}\right) + h_E\left(B^{(1)}, B^{(2)}\right)\right).$$

□

We introduce some notation in order to state the irrepresentability condition of Lee et al. (2015). For a set $F \subset \mathbb{R}^{N\times 2N}$ and $Y \in \mathbb{R}^{N\times 2N}$, denote by $\gamma_F(Y) = \inf\{\lambda > 0 : Y \in F\}$ the gauge function on C. Thus,

$$\gamma_I\left(B^{(1)}, B^{(2)}\right) = \max\left\{\max_{i\in\mathcal{G}^C} \|B^{(1)}_{(i)}\|_2, \ \frac{1}{\rho}\max_{(i,j)\in(\mathcal{G}\times\mathcal{G})^C} |B^{(2)}_{ij}|\right\} + \mathbf{1}_I\left(B^{(1)}, B^{(2)}\right),$$

where $\mathbf{1}_I(B) = 0$ if $B \in I$ and ∞ otherwise. Define

$$V(Z) = \inf\{\gamma_I(Y) : \quad Z - Y \in E^{\perp}, Y \in \mathbb{R}^{N\times 2N}\}$$

for $Z \in \mathbb{R}^{N\times 2N}$. Let $\tilde{\mathcal{M}} = E \cap \operatorname{span}(I)^{\perp}$ be the set of matrices with correct support in the extended space $\mathbb{R}^{N\times 2N}$, similarly to $\mathcal{M}$ in (4.1). Denote by P_M and $P_{M^{\perp}}$ the projections onto $\tilde{\mathcal{M}}$ and $\tilde{\mathcal{M}}^{\perp}$. Define the function $\mathcal{H}(Z) : \mathbb{R}^{N\times N} \to \mathbb{R}^{N\times N}$ as

$$\mathcal{H}(Z)_{ij} = \begin{cases} \operatorname{Tr}\left(H_{(i,j),\mathcal{G}}(P_M Z)_{\mathcal{G},\mathcal{G}}\right) & \text{if } j \in \mathcal{G}, \\ 0 & \text{otherwise.} \end{cases}$$

where $H_{(i,j),\mathcal{G}}$ is the matrix defined in (4.2). The Irrepresentability Assumption 3.2 of Lee et al. (2015) requires the existence of $0 < \tilde{\tau} < 1$ such that

$$\sup_{Z\in A} V\left(P_{M^{\perp}}\left(R\mathcal{H}(Z) - Z\right)\right) < 1 - \tilde{\tau}. \tag{A.2}$$

For a support function h, denote by $\partial h(M) = \bigcup_{Y\in M} \partial h(Y)$ the set of subdifferentials of h in M. Note that $\partial h_A(M) = A$, since $0 \in M$ and $\partial h_A(0) = A$.

Lemma 2. *If Assumption 2 holds and $\rho > \frac{1}{\tau} - \frac{1}{\sqrt{G}}$, then there exists $0 < \tilde{\tau} < 1$ such that* (A.2) *holds.*

Proof of Lemma 2. Since V is sublinear (Lemma 3.3 of Lee et al. (2015)),

$$\sup_{Z\in A} V\left(P_{M^{\perp}}\left(R\mathcal{H}(Z) - Z\right)\right) \le \sup_{Z\in A} V\left(P_{M^{\perp}}\left(R\mathcal{H}(Z)\right)\right) + \sup_{Z\in A} V\left(P_{M^{\perp}} Z\right). \tag{A.3}$$

To bound the first term, note that $E^{\perp} = \{(Z^{(1)}, Z^{(2)}) | Z^{(1)}_{ij} + Z^{(1)}_{ji} + Z^{(2)}_{ij} = 0, 1 \le j < i \le N\}$.

$$\begin{aligned} V(Y^{(1)}, Y^{(2)}) &= \inf\left\{\gamma\left(U^{(1)}, U^{(2)}\right) :\ (U^{(1)}_{ij} - Y^{(1)}_{ij}) + (U^{(1)}_{ji} - Y^{(1)}_{ji}) + (U^{(2)}_{ij} - Y^{(2)}_{ij}) = 0,\ 1 \le j < i \le N\right\} \\ &\le \inf\left\{\gamma(U^{(1)}, U^{(2)}) :\ U^{(1)}_{(i)} = Y^{(1)}_{(i)}, i \in \mathcal{G}^C; U^{(2)}_{\mathcal{G}^C,\mathcal{G}^C} = Y^{(2)}_{\mathcal{G}^C,\mathcal{G}^C};\ (U^{(1)}, U^{(2)}) - (Y^{(1)}, Y^{(2)}) \in E^{\perp}\right\} \\ &\le \max\left\{\max_{i\in\mathcal{G}^C} \|Y^{(1)}_{(i)}\|_2,\ \frac{1}{\rho}\left\|Y^{(2)}_{\mathcal{G}^C,\mathcal{G}^C}\right\|_{\infty}\right\}. \end{aligned}$$

Therefore,

$$\begin{aligned} V\left(P_{M^{\perp}}\left(R\mathcal{H}(Z)\right)\right) &\le \max\left\{\max_{i\in\mathcal{G}^C}\left\|\left(P_{M^{\perp}}\left(R\mathcal{H}(Z)\right)\right)^{(1)}_{(i)}\right\|_2,\ \frac{1}{\rho}\left\|\left(P_{M^{\perp}}\left(R\mathcal{H}(Z)\right)\right)^{(2)}_{\mathcal{G}^C,\mathcal{G}^C}\right\|_{\infty}\right\} \\ &= \max_{i\in\mathcal{G}^C}\left\|\mathcal{H}(Z)_{(i)}\right\|_2, \end{aligned}$$

which implies that

$$\begin{aligned} \sup_{Z\in A} V\left(P_{M^{\perp}}\left(R\mathcal{H}(Z)\right)\right) &\le \sup_{Z\in A}\left\{\max_{i\in\mathcal{G}^C}\left\|\mathcal{H}(Z)_{(i)}\right\|_2\right\} \\ &\le \sup_{B\in\mathbb{R}^{G\times G}, \|B_{(i)}\|_2\le 1}\left\{\max_{i\in\mathcal{G}^C}\left\|\left(\operatorname{Tr}\left(H_{(i,j),\mathcal{G}}B\right)\right)_{j=1}^{N}\right\|_2\right\} \\ &\le \max_{i\in\mathcal{G}^C}\left\|\left(\sum_{k=1}^{G}\|(H_{(i,j),\mathcal{G}})_{k\cdot}\|_2\right)_{j=1}^{N}\right\|_2 = 1 - \tau. \end{aligned} \tag{A.4}$$

Let $Z = (Z^{(1)}, Z^{(2)}) \in A$. Without loss of generality, assume that $Z^{(1)}_{\mathcal{G},\mathcal{G}} = Z^{(2)}_{\mathcal{G},\mathcal{G}} = 0$ (note that these entries

do not change $V(P_{M^\perp}Z)$). Therefore, $P_{M^\perp}Z = Z$. Hence,

$$\begin{aligned}
V(Z) &= \inf\left\{\gamma\left(U^{(1)},U^{(2)}\right): \ \left(U^{(1)},U^{(2)}\right)\in I,\ \left(U^{(1)},U^{(2)}\right)-\left(Z^{(1)},Z^{(2)}\right)\in E\right\}\\
&= \inf\left\{\gamma\left(U^{(1)},U^{(2)}\right): \ U^{(1)}_{ij}+U^{(2)}_{ij}=Z^{(1)}_{ji}, 1\le j\le G, G<i\le N\right\}\\
&= \inf\left\{\max\left\{\max_{i\in\mathcal{G}^C}\|U^{(1)}_{(i)}\|_2, \frac{1}{\rho}\max_{\substack{1\le j\le G\\ G<i\le N}}|U^{(2)}_{ij}|\right\}: \ U^{(1)}_{ij}+U^{(2)}_{ij}=Z^{(1)}_{ji}, 1\le j\le G, G<i\le N\right\}\\
&\le \inf\left\{\max\left\{\max_{i\in\mathcal{G}^C}\|U^{(1)}_{(i)}\|_2, \frac{1}{\rho}\max_{\substack{1\le j\le G\\ G<i\le N}}|U^{(2)}_{ij}|\right\}: \ U^{(1)}_{ij}+U^{(2)}_{ij}=1, 1\le j\le G, G<i\le N\right\}
\end{aligned}$$

The last bound from $|Z^{(1)}_{ji}|\le 1$ and no longer depends on Z. It is easy to see that the minimum is attained when, for each $i>G$,

$$\left\|U^{(1)}_{(i)}\right\|_2 = \frac{1}{\rho}\left|U^{(2)}_{ij}\right|, \quad 1\le j\le G,$$

and therefore

$$V(Z)\le \frac{\sqrt{G}}{1+\rho\sqrt{G}}. \tag{A.5}$$

Moreover, if $Z^*\in A$ is defined such that $(Z^*)^{(1)}_{G+1,i}=1$ for $i=1,\dots,G$ and 0 elsewhere, then $V(Z^*)$ achieves this bound, which shows that $\rho>1-\frac{1}{\sqrt{G}}$ is a necessary condition for the irrepresentability to hold, even in the case where the entries of the Hessian that denote the information between active and inactive edges is zero. Therefore, plugging the bounds (A.4) and (A.5) into equation (A.3), we obtain (A.2) holds as long as $1-\tau+\frac{\sqrt{G}}{1+\rho\sqrt{G}}<1$, which implies that $\rho>\frac{1}{\tau}-\frac{1}{\sqrt{G}}$. □

The next lemma establishes a bound on the dual norm of Ω of the loss gradient function under a sub-Gaussian assumption. Let Ω^* denote the dual norm of Ω, so $\Omega^*(B)=\sup\{\langle Y,B\rangle \mid Y\in C,\Omega(Y)\le 1\}$.

Lemma 3. *Under Assumption 3,*

$$\mathbb{P}\left(\Omega^*(\nabla\ell(B^\star))>t\right)\le 2\exp\left(2\log N-\frac{n(1+\rho)^2t^2}{2N(\sigma^2}\right).$$

Proof of Lemma 3. By Hoeffding's inequality for sub-Gaussian variables, for all j,k and $t>0$,

$$\mathbb{P}\left(|\nabla_{jk}\ell\left(B^\star\right)|>t\right)\le\mathbb{P}\left(\left|\frac{1}{n}\sum_{i=1}^n\nabla_{jk}\ell_i\left(B^\star\right)\right|>t\right)\le 2\exp\left(-n\frac{t^2}{\sigma^2}\right).$$

Note that $(1+\rho)\sum_{i=1}^N\|B_{(i)}\|_2\le\Omega(B)$. Let $\Phi(B)=\frac{1}{1+\rho}\max_{i=1,\dots N}\|B_{(i)}\|_2$. Thus,

$$\Omega^*(B)\le\sup_{Y\in\mathbb{R}^{N\times N}}\{\mathrm{Tr}(YB):\Omega_{\rho=0}(Y)\le 1\}=\Phi(B). \tag{A.6}$$

Setting $\Lambda = \nabla \ell(B^\star)$,

$$\begin{aligned}
\mathbb{P}\left(\Omega^*(\Lambda) > t\right) &\leq \mathbb{P}\left(\Phi(\Lambda) > t\right) \\
&= \mathbb{P}\left(\max_{1\leq i\leq N} \|\Lambda_{(i)}\|_2 > (1+\rho)t\right) \\
&\leq \mathbb{P}\left(\max_{1\leq i\leq N} \max_{j\neq i} |\Lambda_{ij}| > (1+\rho)\frac{t}{\sqrt{N}}\right) \\
&\leq 2N(N-1)\exp\left(-\frac{n(1+\rho)^2 t^2}{2\sigma^2(N-1)}\right),
\end{aligned}$$

the last inequality obtained by arguments similar to Lemma 4.3 of Lee et al. (2015). □

For a semi-norm $\Psi : \mathbb{R}^{N\times N} \to \mathbb{R}$, let κ_Ω be the compatibility constant between Ψ and the Frobenius norm, defined as

$$\kappa_\Psi = \sup\left\{\Psi(B) : \|B\|_2 \leq 1, B \in M\right\},$$

and let κ_{IC} be the compatibility constant between the irrepresentable term and the dual norm Ω^* given by

$$\kappa_{\mathrm{IC}} = \sup\left\{V(P_{M^\perp}(R\mathcal{H}Z - Z)) : \Omega^*(Z) \leq 1\right\} \ .$$

Lemma 4. *The following bounds on the compatibility constants hold:*

$$\begin{aligned}
\kappa_\Omega &= \sqrt{G} + \rho\sqrt{G(G-1)}, \\
\kappa_{\Omega^*} &\leq \frac{1}{1+\rho}, \\
\kappa_{IC} &\leq 3 - \tau.
\end{aligned}$$

Proof of Lemma 4. Note that $\Omega(Y)$ is maximized on $\{Y : \|Y\|_2 \leq 1\}$ when all entries of Y have magnitude equal to $\frac{1}{\sqrt{G(G-1)}}$. Therefore

$$\kappa_\Omega = G\sqrt{\frac{G-1}{G(G-1)}} + \rho\frac{G(G-1)}{\sqrt{G(G-1)}} = \sqrt{G} + \rho\sqrt{G(G-1)}. \tag{A.7}$$

Similarly, (A.6) implies

$$\kappa_{\Omega^*} \leq \sup\left\{\frac{1}{1+\rho}\max_{i\in\mathcal{G}} \|B_{(i)}\|_2 : \|B\|_2 \leq 1\right\} \leq \frac{1}{1+\rho}. \tag{A.8}$$

Finally,

$$\begin{aligned}
\kappa_{IC} &= \sup\left\{V(P_{M^\perp}(R\mathcal{H}Z - Z)) : \Omega^*(Z) \leq 1\right\} \\
&\leq \sup\left\{V(P_{M^\perp}(R\mathcal{H}Z)) : \Omega^*(Z) \leq 1\right\} + \sup\left\{V(P_{M^\perp}(Z)) : \Omega^*(Z) \leq 1\right\} \\
&\leq (1-\tau) + 2 = 3 - \tau.
\end{aligned}$$

□

Proof of Proposition 1. Part (a). Since $\hat{B}$ minimizes the problem (2.4),

$$\ell(\hat{B}) + \lambda\Omega(\hat{B}) \leq \ell(B^\star) + \lambda\Omega(B^\star).$$

Rearranging the terms, using Assumption 1, by the triangle inequality and the generalized Cauchy-Schwarz inequality,

$$\begin{aligned} 0 &\geq \ell(\hat{B}) - \ell(B^\star) + \Omega(\hat{B}) - \Omega(B^\star) \\ &\geq \left\langle \nabla\ell(B^\star)^T, \hat{B} - B^\star \right\rangle + \frac{m}{2}\|\hat{B} - B^\star\|_2^2 - \Omega\left(\hat{B} - B^\star\right) \\ &\geq -\Omega\left(\hat{B} - B^\star\right)\Omega^*\left(\nabla\ell(B^\star)\right) - \Omega\left(\hat{B} - B^\star\right) + \frac{m}{2}\|\hat{B} - B^\star\|_2^2. \end{aligned} \tag{A.9}$$

Using the argument for computing κ_Ω in (A.7), $\Omega(Y) \leq \left(\sqrt{N} + \rho\sqrt{N(N-1)}\right)\|Y\|_2$. Rearranging the terms in (A.9),

$$\|\hat{B} - B^\star\|_2 \leq \frac{2}{m}\left(\sqrt{N} + \rho\sqrt{N(N-1)}\right)\left(\lambda + \Omega^*\left(\hat{B} - B^\star\right)\right).$$

Setting $\lambda = K\sqrt{\frac{3N\log N}{n}}$, by Lemma 3, with probability at least $1 - 2/N$,

$$\|\hat{B} - B^\star\|_2 \leq \frac{4}{m}\left(\sqrt{N} + \rho\sqrt{N(N-1)}\right)\lambda.$$

Part (b). Lemma 1 gives a geometric decomposition of the penalty. Therefore, we can directly use Theorem 3.1 of Lee et al. (2015), since Lemma 2 also ensures that their irrepresentability condition holds. Thus,

$$\|\hat{B} - B^\star\|_2 \leq \frac{2}{m}\kappa_\Omega\left(1 + \frac{\tau}{4\kappa_{\mathrm{IC}}}\right)\lambda,$$

and $\hat{\mathcal{G}} \subseteq \mathcal{G}$ as long as

$$\frac{4\kappa_{\mathrm{IC}}}{\tau}\Omega^*\left(\nabla\ell(B^\star)\right) < \lambda < \frac{m^2\tau}{2L\kappa_\Omega^2\kappa_{\Omega^*}\kappa_{\mathrm{IC}}}\left(1 + \frac{\tau}{4\kappa_{\mathrm{IC}}}\right)^{-2}.$$

Setting

$$\lambda = \frac{4K\kappa_{\mathrm{IC}}}{\tau}\sqrt{\frac{3N\log N}{n}},$$

the first inequality holds with probability at least $1 - 2/N$. The second inequality holds as long as the sample size requirement is valid. Therefore, claims (4.5) and (4.6) follow.

□